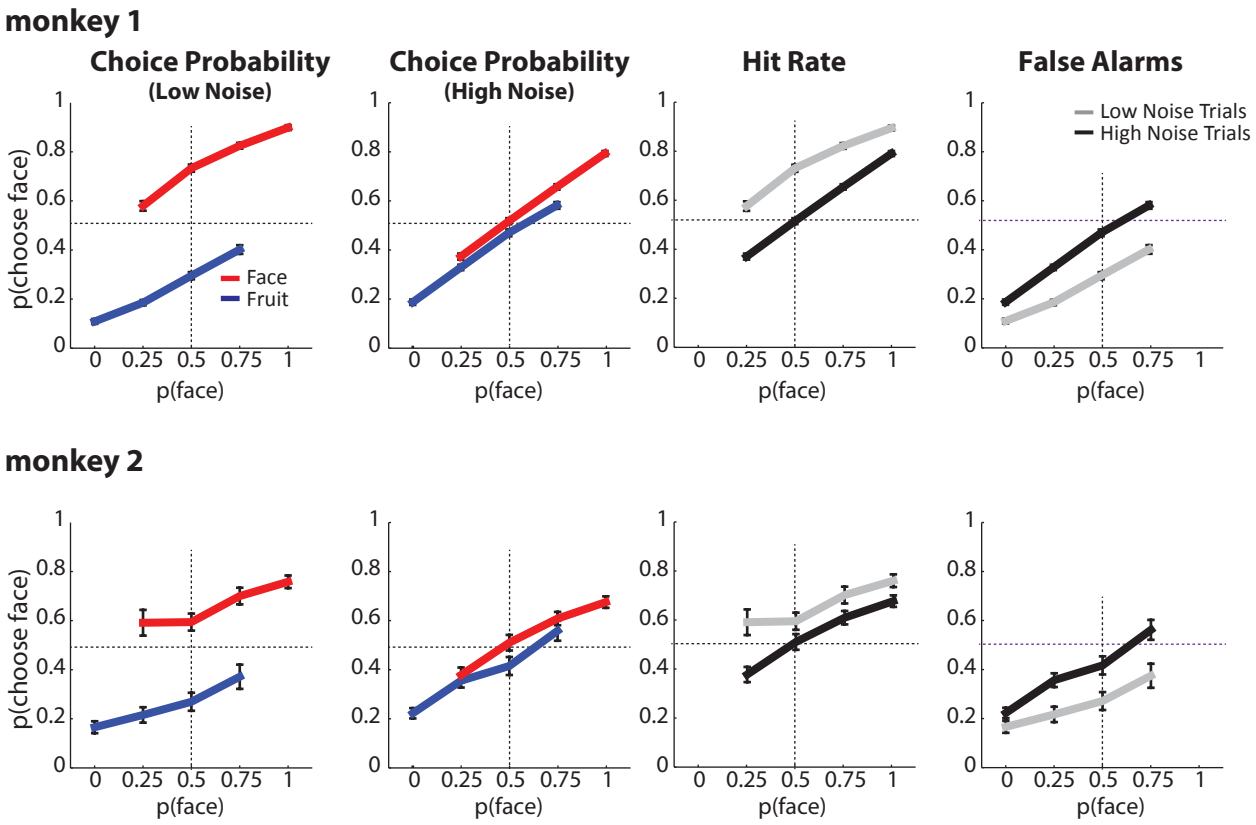

Supplemental Figure 1

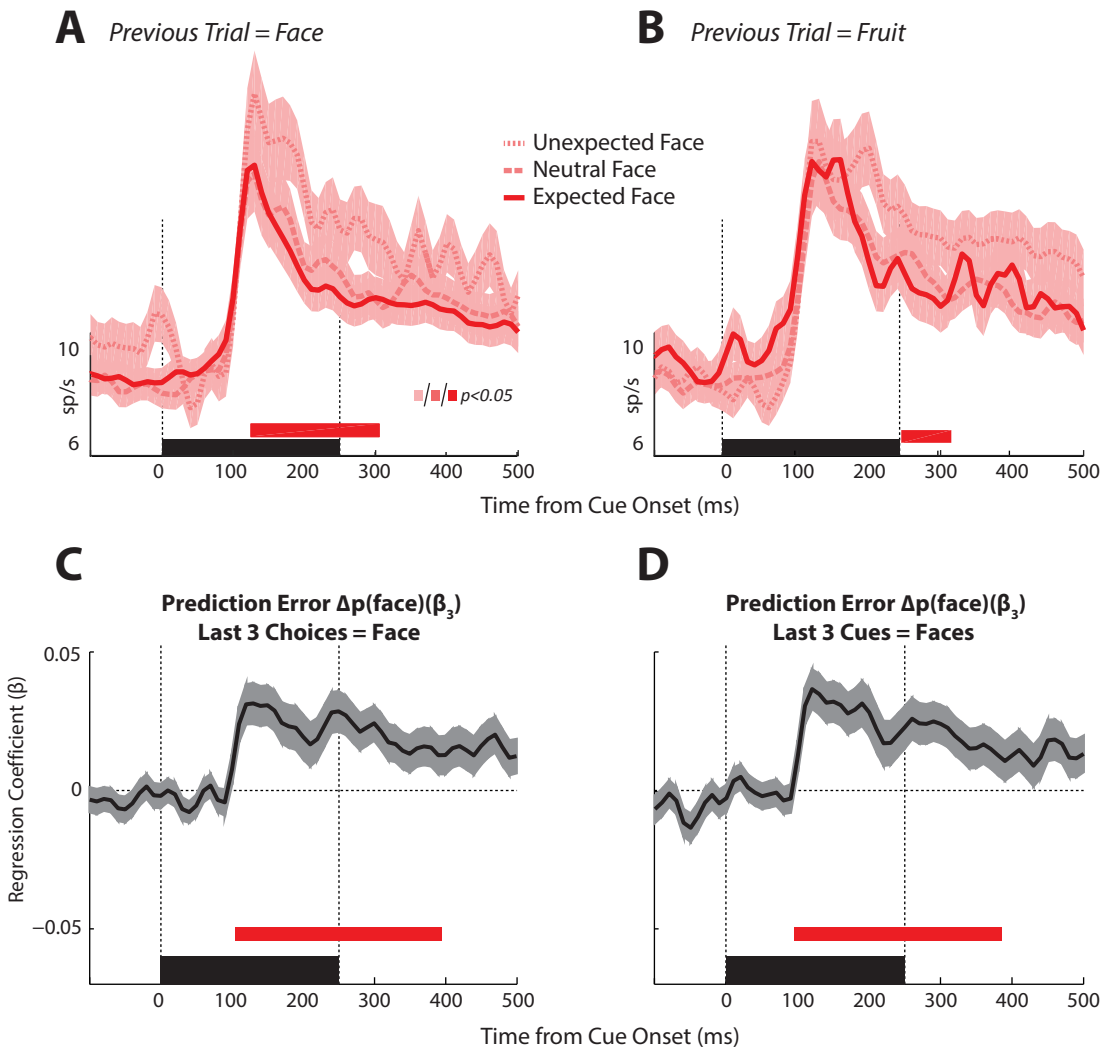

Supplemental Figure 2

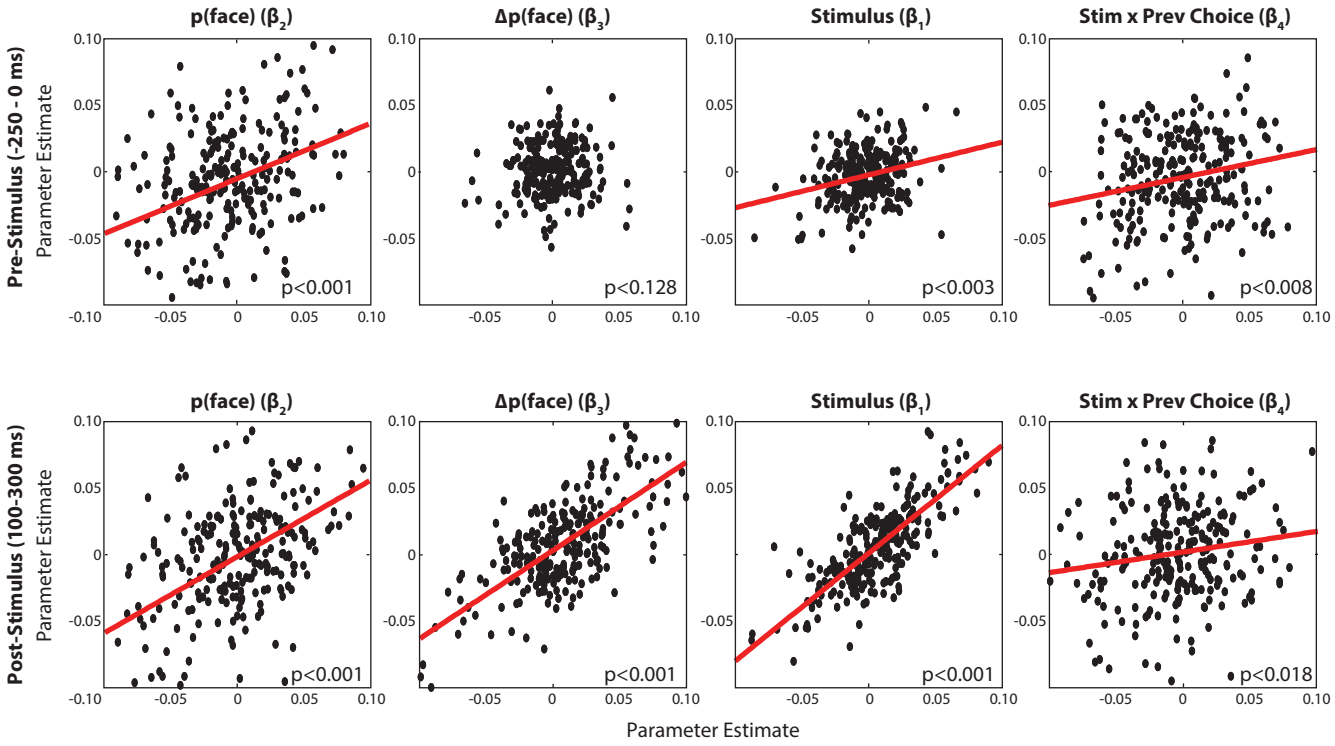

Supplemental Figure 3

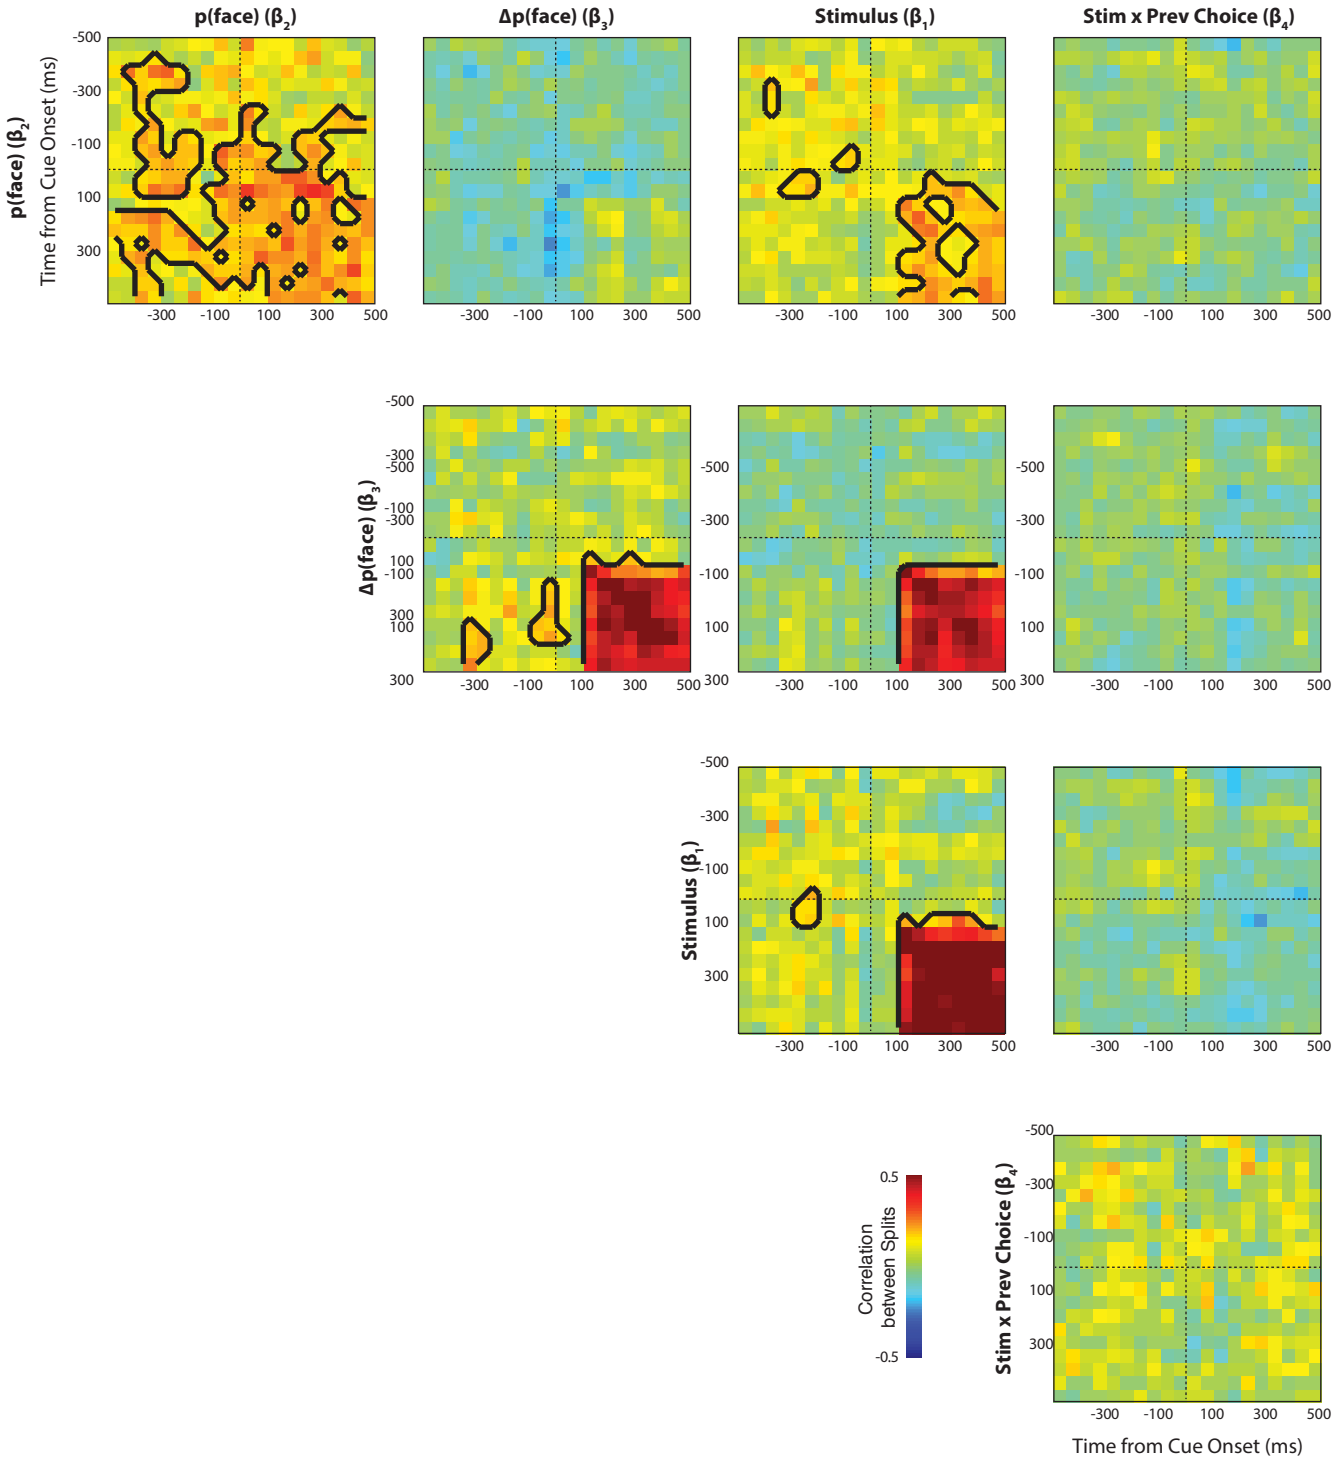

Supplemental Figure 4

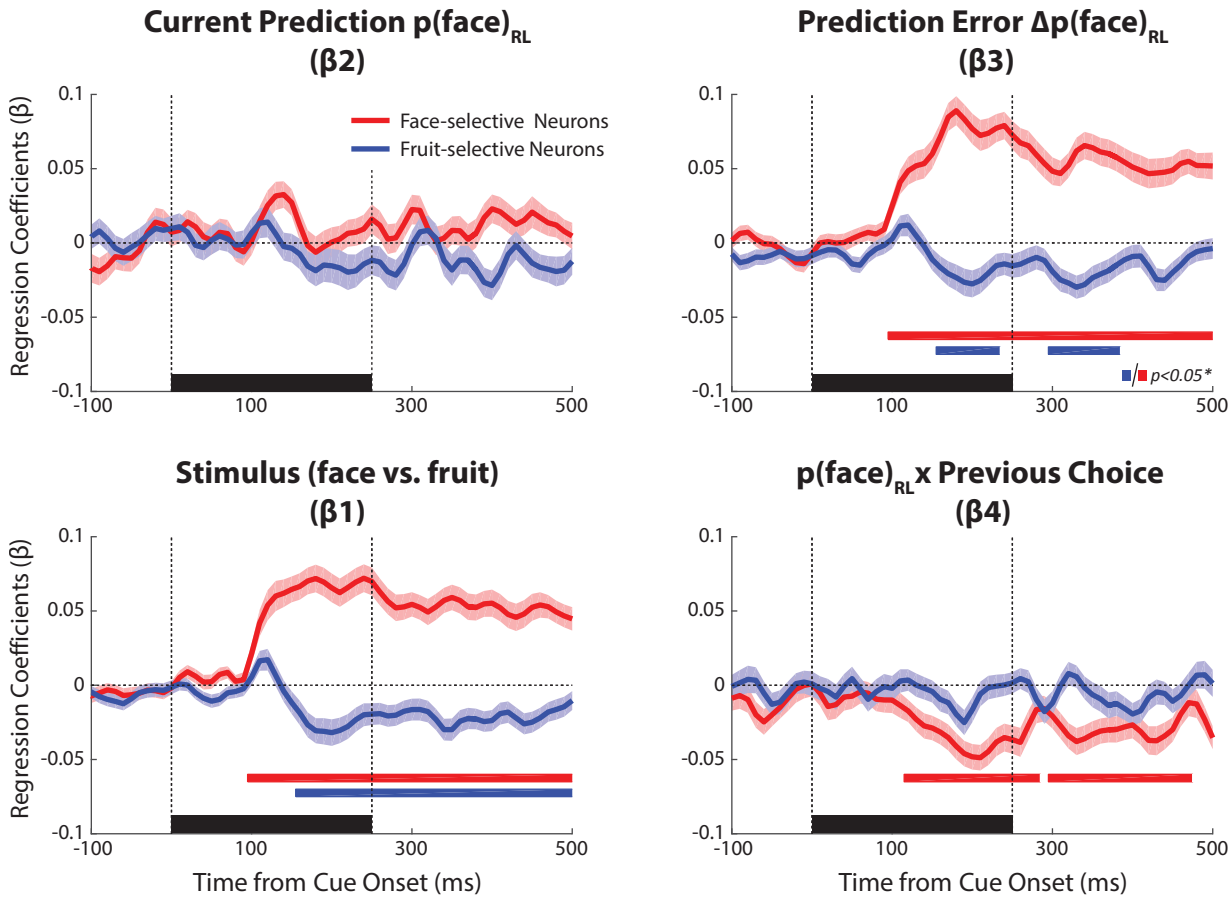

Supplemental Figure 5

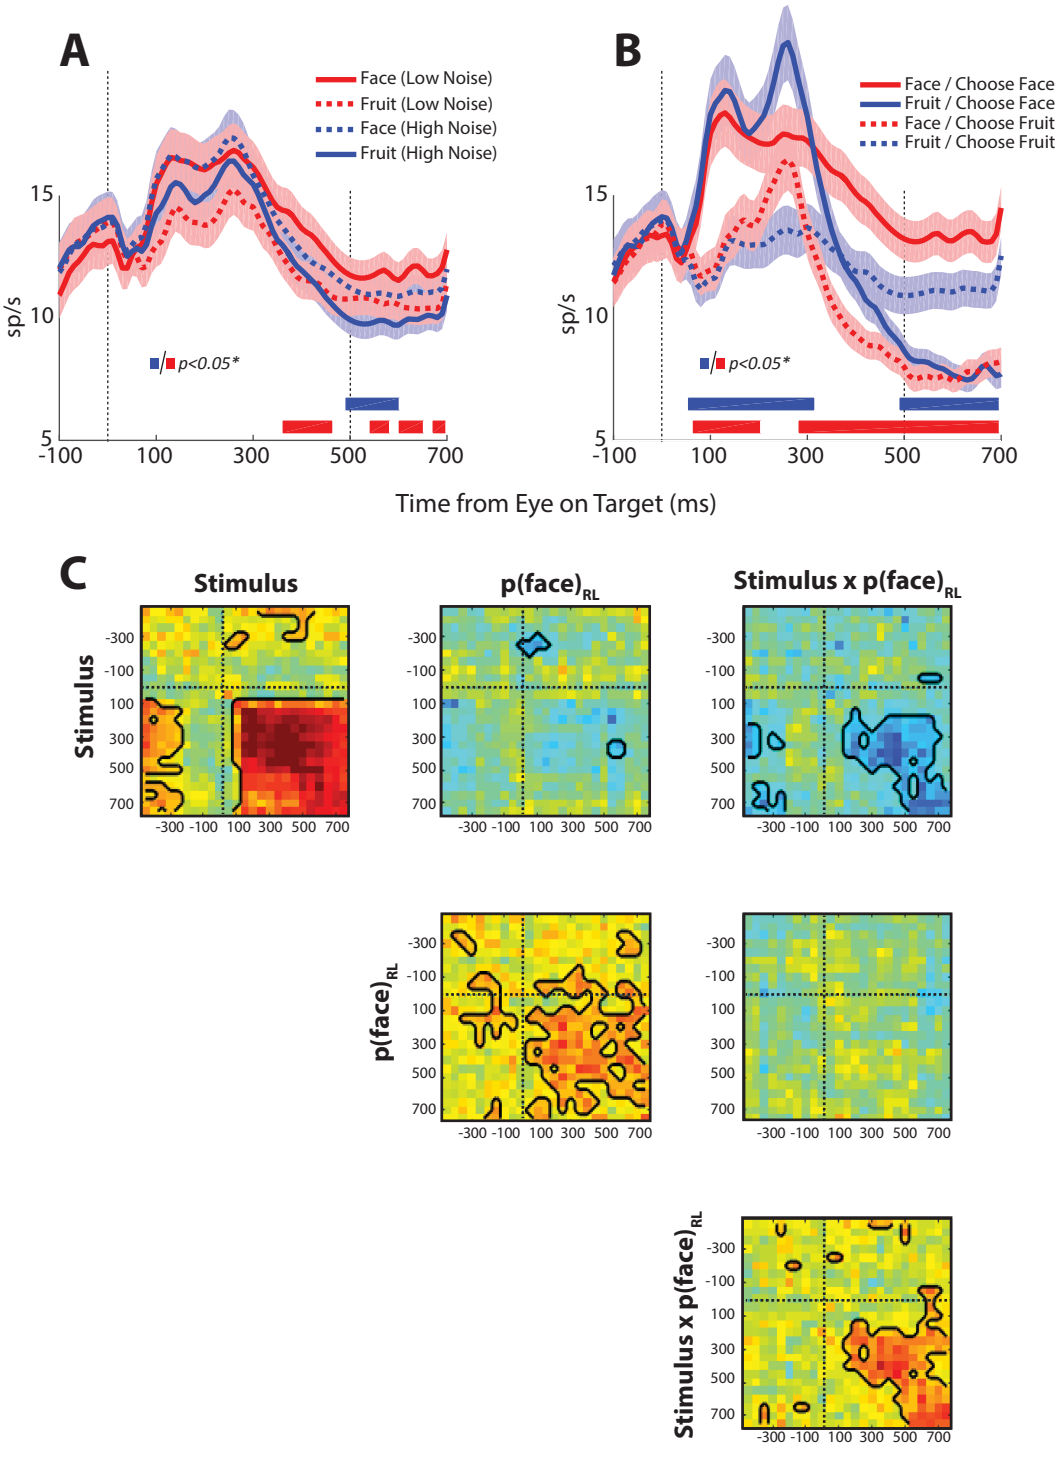

Supplemental Figure 6

Supplemental Figures Legends

Figure S1, related to Figure 1 Behavioural performance in delayed match-to-sample task for each monkey. Left Panels: choice probabilities (for choosing face) in response to low and high noise cues. Right Panels: Hit rates and false alarms were more influenced by *p(face)* under high noise conditions as compared to low noise conditions (grey lines). SEM indicated by error bars.

Figure S2, related to Figure 3 Population **r**esponses in IT to expected, neutral, and unexpected faces, grouped according to the identity of the cue stimulus from the previous trial. Regardless of the nature of the previous trial, a robust and sustained increase in activity in response to face stimuli was observed, beginning around 100 ms following cue onset. Black bar indicates cue presentation period. Red bars indicate timepoints where average firing rates for expected and unexpected faces were significantly different from one another. (C,D) Cue-aligned regression coefficients for the interaction between stimulus category and p(face) ("prediction error", β3) for Eq.1(which includes a regressor based on previous choice) and for the following equation, which includes a regressor for the three previous cue stimuli (prevstim):

$$y = \beta_0 + \beta_1 stimulus + \beta_2 p(face)_{Bayes} + \beta_3 stimulus \times p(face)_{Bayes} + \beta_4 prevstim_1 \times stimulus + \beta_5 prevstim_2 \times stimulus + \beta_6 prevstim_3 \times stimulus + \beta_7 trial$$

The red bars show the corresponding significant timepoints for each analysis. To correct for multiple comparisons, a cluster-correction was applied across timepoints ($p<0.05$*).

Figure S3, related to Figure 4 Scatter plots of the parameter estimates for all coefficients in eq. 4, calculated from the two independent splits of the data, both in the pre-stimulus period (top row) and post-stimulus period (bottom row). Inset text shows the p-value for the corrected non-parametric test (see Experimental Procedures).

Figure S4, related to Figure 5 Cross-validation of all choice-aligned regression coefficients. Each plot shows the Pearson's correlation between regression coefficients for two independent splits of the data. Panels on the diagonal axis show correlations for the same predictor variable (e.g., β_2 and β_2; top row, leftmost panel) whereas off-diagonal panels show correlations for between predictor variables (e.g., β_2 and β_4, top row, rightmost panel). Each panel shows correlations for a given timepoint with every other timepoint. Black contour lines indicate timepoints where a significant correlation was observed ($p < 0.01$), adjusted for multiple comparisons using a cluster correction method.

Figure S5, related to Figure 5 Cue-aligned regression coefficients from eq. 4 for face and fruit. Black bar indicates cue presentation period. Red and blue traces correspond to separate regressions for different subpopulations of neurons, sorted by their preference for faces (red) and fruit (blue) based on a median split on the regression coefficients for a simple regression of stimulus category on neural activity. The red and blue bars show the corresponding significant timepoints for each analysis. To correct for multiple comparisons, a cluster-correction was applied across timepoints ($p<0.05$*).

Figure S6, related to Figure 6 Average responses and regression coefficients for IT neurons at the time of choice. Responses are aligned on the time when the monkey's eye reaches the chosen target stimulus (either an intact face or fruit). The monkey is then required to hold fixation on the target stimulus for 500 ms before a liquid reward is administered. (A) Average response, grouped according to *cue identity.* Red and blue bars indicate timepoints where average responses on high vs. low noise trials were significantly different from one another. (B) Average responses, grouped according to correct vs. incorrect trials. Solid traces indicate trials where the monkey chose "face", dashed traces are those trials where the monkey chose "fruit". Red traces indicate trials where the cue was a face, blue traces indicate trials where the cue was a fruit. Red and blue bars indicate timepoints where average responses for correct vs. incorrect trials were significantly different from one another. To correct for multiple comparisons, a cluster-correction was applied across timepoints ($p<0.05$*). (C) Cross-validation of choice-aligned regression coefficients. Each plot shows the Pearson's correlation between regression coefficients for two independent splits of the data. Panels on the diagonal axis show correlations for the same predictor variable (e.g., β_1 and β_1; top row, leftmost panel) whereas off-diagonal panels show correlations for between predictor variables (e.g., β_1 and β_2, top row, middle panel). Each panel shows correlations for a given timepoint with every other timepoint. Black contour lines indicate timepoints where a significant correlation was observed ($p < 0.01$), adjusted for multiple comparisons using a cluster correction method.

Supplemental Experimental Procedures

All procedures were approved by the National Institute of Mental Health (NIMH) Animal Care and Use Committee and conformed to all NIH guidelines. Two adult, male rhesus macaques (10-12 kg) were prepared for chronic recordings. A post for head restraint and recording cylinders were implanted under aseptic conditions. Circular recording chambers (Crist Instruments, Hagerstown MD) were centred about 12 mm anterior to the inter-aural axis over the right hemisphere in both animals.

Monkeys were trained on a delayed match-to-sample task that required them to identify which of two canonical stimuli (a face or fruit) best matched a previously presented cue (**Fig. 1A**). Each trial began with an initial fixation period of 400-800 ms, after which a cue stimulus was presented for 250 ms. The cue stimulus was either a degraded face or a fruit (face and fruit exemplars were randomly selected on a daily basis from a set of 8 from each category). Stimulus degradation was accomplished through the addition of Gaussian noise (Adobe Photoshop, Adobe Systems, San Jose, CA) at two different levels (20%, "low noise" or 80%, "high noise"). The cue was removed and after an additional delay period (randomised between 150-550 ms), intact face and fruit stimuli (where one matched the identity of the degraded cue stimulus) appeared on the right and left side of the fixation point. The monkey was required to generate a saccade to the matching stimulus and hold fixation for 500 ms before receiving a fluid reward. The next trial began after a 500-1000 ms (randomised) inter-trial interval during which time no stimulus was present and the monkeys were free to move their eyes. The cue and choice stimuli were converted to grayscale and vignetted with an oval window to remove low-level cues such as overall shape. They were approximately 5 degrees in size, and the choice stimuli were positioned about 10 degrees on either side of the fixation point. The respective side on which the canonical face/fruit stimuli appeared was pseudorandomised.

Trials were arranged into 5 blocks, (presented in random order), each with a different probability of the cue being a face vs. a fruit (0%, 25%, 50%, 75%, 100%). The monkeys were given no external cue as to which block they were currently in, nor when blocks transitioned from one probability to the next. The monkeys had to perform at least 3-5 correct trials per condition (approximately 50 trials per block).

Neuronal data were collected and processed using methods described in detail in a previous paper [S1]. Data were collected over 119 recording sessions (monkey 1: 97 sessions, 217 neurons; monkey 2: 22 sessions, 36 neurons). For welfare reasons, we were required to end data collection prematurely in monkey 2, hence the reduced number of recording sessions and neurons. However, no significant differences were observed between the data obtained from monkey 2 as compared to monkey S and so we grouped data across the two animals.

During recording sessions, between 1 and 4 electrodes were lowered into the inferior bank and lip of the superior temporal sulcus (between 5 and 19 mm anterior to the interaural axis), guided by rigid guidetubes that terminated about 10 mm above the targeted area. Waveform data was sampled at 40KHz and later sorted into individual neurons using Offline-Sorter (Plexon Systems). Spiking data was convolved with a Gaussian kernel (σ = 10 ms) to generate spike density functions for each trial.

Behaviour and modelling. Data from both monkeys were pooled as if they were a single observer. Behavioural data were analysed using analysis of variance (ANOVAs). For initial behavioural analyses, we used the objective task structure, measuring the probability of responding "face" as a function of the stimulus (face vs. fruit), the noise level (high vs. low), and the objective probability of a face, *p(face)*. Subsequently, we fit a Bayesian Learner [S2] to the data. Briefly, the Bayesian learner calculates a posterior likelihood distribution over *p(face)* and V, the rate of change of *p(face)*, following each new stimulus, marginalising over V to obtain $p(face)_{Bayes}$ (code adapted from: http://hannekedenouden.ruhosting.nl/RLtutorial/Instructions.html). The delta-rule RL model was implemented as described in the main text. The predictions of delta-rule and Bayesian Learner models were compared using maximum likelihood estimation [S3], calculating single-trial likelihoods p(data|model) by combining psychophysical data (probability that the monkey made a face response in that condition) with model-derived estimates of *p(face)*. Model probabilities and exceedance probabilities were calculated via Bayesian model selection [S4].

Neuronal data. Firing rates on each trial were averaged over bins of 10 ms. Average data were plotted for each trial type with standard errors computed over all 253 neurons. To assess neuronal selectivity, we compared baseline (-250 to 0 ms) to post-stimulus (100 to 300 ms) firing rates for low-noise stimuli only using Wilcoxon signed rank tests for each of the 253 neurons. We used this conventional approach for comparability with previous studies, but very similar results were obtained when using regression-based methods for defining selectivity (e.g., **Fig. 3F**).

Multiple regression was performed for each neuron separately to estimate beta coefficients associated with predictor variables, as described in the main text. Significance was calculated by performing t-tests on the resulting coefficients at each timepoint, using an alpha of $p < 0.05$. To correct for multiple comparisons, we used a cluster permutation approach [S5]. The trial structure within a neuron was shuffled 100 times, and the analysis was repeated (preserving all temporal aspects of the data). For each shuffle, the maximum number of adjacent significant timepoints (cluster size) was logged. Only clusters falling within the 95th percentile of the resulting distribution are reported.

Multivariate analyses were conducted by splitting the data from each neuron into two random sets of trials of equal number, and estimating the coefficients from the regression model separately for each half of the data. The two sets of coefficients were then correlated between independent splits for each timepoint. Significance was calculated by comparing nonparametric (i.e. rank-based) Fisher's Z-scores obtained from this correlation to a null distribution obtained by shuffling the data 1000 times and repeating the analysis. Timepoints where the correlation exceeded $p < 0.05$ (corrected for multiple comparisons) are marked with contours. The same cluster permutation method based on shuffled data was again used to exclude smaller clusters, except that clusters were defined by a surface connectivity criterion [S6].

Multivariate decoding. Neuronal data were first sorted so that trials from each condition (e.g., neutral face stimulus, neutral fruit stimulus) fell in an equal number of adjacent columns, with excess trials discarded. Critically, this ensured equal numbers of trials per condition, ensuring that our analyses were not biased by the larger number of trials in the "expected" condition. Neurons with less than 20 trials in a condition were excluded (~20%). Data were then randomly allocated to training (70%) and test (30%) datasets. Using probit regression, coefficients predicting face stimulus vs. fruit stimulus were estimated from the training set for each timepoint, and used to predict stimulus classes for the corresponding timepoint on the test set, according to whether the resulting predicted value exceeded 0.5. This was carried out separately for neutral trials (where $0.33 < p(face)\text{Bayes} < 0.66$) and expectation trials (where the expectation and stimulus were congruent: i.e., pooling over cases were $p(face)_{\text{Bayes}} > 0.66$ and the stimulus was face, or $p(face)_{\text{Bayes}} < 0.33$ and the stimulus was fruit. This entire process was performed separately for high noise and low noise trials. The resulting decoding accuracy was plotted over time in each condition, in both the cue and response periods.

References

[S1] Bell, A.H., Malecek, N.J., Morin, E.L., Hadj-Bouziane, F., Tootell, R.B., and Ungerleider, L.G. (2011). Relationship between functional magnetic resonance imaging-identified regions and neuronal category selectivity. J Neurosci *31*, 12229-12240.

[S2] Behrens, T.E., Woolrich, M.W., Walton, M.E., and Rushworth, M.F. (2007). Learning the value of information in an uncertain world. Nat Neurosci *10*, 1214-1221.

[S3] Daw, N. (2011). Trial-by-trial data analysis using computational models. In Decision Making, Affect, and Learning: Attention and Performance XXIII, M. Delgado, E.A. Phelps and T.W. Robbins, eds. (Oxford, UK).

[S4] Stephan, K.E., Penny, W.D., Daunizeau, J., Moran, R.J., and Friston, K.J. (2009). Bayesian model selection for group studies. Neuroimage *46*, 1004-1017.

[S5] Maris, E., and Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. Journal of neuroscience methods *164*, 177-190.

[S6] Kersten, D., Mamassian, P., and Yuille, A. (2004). Object perception as Bayesian inference. Annu Rev Psychol *55*, 271-304.